ETH Zürich
Department of Physics

Master Thesis

Implementing spectral methods for hidden Markov models with real-valued emissions

Author:
Carl Mattfeld[1,2]

Supervisors:
Prof. Dr. Chris Wiggins[1]
Prof. Dr. Didier Sornette[2]

[1] Columbia University, New York, NY 10027, USA
[2] ETH Zürich, 8092 Zürich, Switzerland

April 30, 2014

arXiv:1404.7472v1 [cs.LG] 29 Apr 2014

To my family.

Acknowledgements

I would like to thank Professor Chris Wiggins for supervising this thesis and challenging me with this project. I am grateful to have had this opportunity to join him as a visiting researcher at Columbia University.
I would like to thank Professor Didier Sornette for being available as my supervisor at ETH Zurich.
Furthermore, I would like to thank Professor Daniel Hsu for the valuable discussions. These discussions were immensely helpful for understanding his original work and getting important ideas, especially for the binning of real-valued emissions.
I am highly grateful for being sponsored by a stipend of the German National Academic Foundation (*Studienstiftung des deutschen Volkes*). This stipend included financial sponsorship for my regular studies in Switzerland, a semester abroad in Hong Kong, and this research project in New York. I would like to thank Professor Dagmar Iber, my liaison professor with the Studienstiftung, and Marius Spiecker gen. Döhmann, my advisor with the Studienstiftung in Bonn.
Finally, I would like to thank my family and my friends for their support and being there for me.

Abstract

Hidden Markov models (HMMs) are widely used statistical models for modeling sequential data. The parameter estimation for HMMs from time series data is an important learning problem. The predominant methods for parameter estimation are based on local search heuristics, most notably the expectation–maximization (EM) algorithm. These methods are prone to local optima and oftentimes suffer from high computational and sample complexity. Recent years saw the emergence of spectral methods for the parameter estimation of HMMs, based on a method of moments approach. Two spectral learning algorithms as proposed by [Hsu *et al.* 2012] and [Anandkumar *et al.* 2012b] are assessed in this work. Using experiments with synthetic data, the algorithms are compared with each other. Furthermore, the spectral methods are compared to the Baum-Welch algorithm, a well-established method applying the EM algorithm to HMMs. The spectral algorithms are found to have a much more favorable computational and sample complexity. Even though the algorithms readily handle high dimensional observation spaces, instability issues are encountered in this regime. In view of learning from real-world experimental data, the representation of real-valued observations for the use in spectral methods is discussed, presenting possible methods to represent data for the use in the learning algorithms.

Contents

1 Introduction **1**
- 1.1 Outline . . . 2
- 1.2 State of research . . . 2
- 1.3 Notation . . . 4
- 1.4 Hidden Markov models . . . 4
  - 1.4.1 Markov process . . . 4
  - 1.4.2 Hidden states as latent variables . . . 5
  - 1.4.3 State transitions . . . 6
  - 1.4.4 Emissions . . . 6
  - 1.4.5 Example . . . 7
  - 1.4.6 Learning problem . . . 8

2 Theory **11**
- 2.1 Spectral algorithm by Hsu et al. . . . 11
  - 2.1.1 Algorithm learnHKZ . . . 12
- 2.2 Spectral algorithm by Anandkumar et al. . . . 13
  - 2.2.1 Bag-of-words . . . 14
  - 2.2.2 Multi-view mixture models . . . 15
  - 2.2.3 Derivations . . . 17
  - 2.2.4 Algorithm learnAHK . . . 19
- 2.3 Real-valued emissions . . . 22
  - 2.3.1 Simple binning approach . . . 22

3 Results and discussion **25**
- 3.1 Comparison of spectral learning algorithms . . . 25
  - 3.1.1 Learning synthetic HMMs . . . 25
  - 3.1.2 Accuracy and dimensionality . . . 27
  - 3.1.3 Comparison with the Baum-Welch algorithm . . . 32
- 3.2 Real-valued emissions . . . 33
  - 3.2.1 Generative process for synthetic experimental data . . . 33
  - 3.2.2 Simple binning approach . . . 33

3.2.3 Alternative approaches 36

4 Conclusion **39**

A Appendix **41**
A.1 Proofs 41
A.2 Specifications of computing environment 43
A.3 Source code 44

List of Figures

1.1 A graphical structure representing a hidden Markov model: a sequence of discrete hidden states $\{h_t\}_{t=1,\dots,L}$ with the corresponding emissions $\{x_t\}_{t=1,\dots,L}$. 5

1.2 A two-state HMM (states 1 and 2) with three possible emissions (1, 2 and 3). The initial state distribution is indicated inside the circles of the corresponding states (π_1 and π_2). Solid arrows indicate transitions between hidden states, dashed arrows indicate emissions of the possible observations. 7

3.1 Comparison of the accuracy and running time of LEARNHKZ and LEARNAHK for $\mathbf{k} = \mathbf{2}$. 28

3.2 Comparison of the accuracy and running time of LEARNHKZ and LEARNAHK for $\mathbf{k} = \mathbf{3}$. 30

3.3 Comparison of the algorithm LEARNAHK (indicated by *AHK*) with the Baum-Welch algorithm, both with a random initialization (*BW*), and using initial values from the LEARNAHK algorithm (*AHK+BW*). Average errors and running times across 100 realizations for *AHK* and 10 realizations for *BW* and *AHK+BW* 31

3.4 The generative process for a sequences of real-valued emissions. The categorical emissions (blue dashed line) were used as the mean values to generate a sequence of normally distributed random numbers with fixed standard deviation (here $\sigma = 0.25$). 34

3.5 Accuracy of the parameter estimates using the simple binning approach. Two-state HMMs with three and six possible emissions, respectively. 35

Chapter 1

Introduction

Hidden Markov models (HMMs) are important statistical models used to describe real-world processes that produce sequential data. The underlying systems are assumed to be Markov processes, with the states of the processes not directly observable (*hidden*). Only a sequence of emissions is observable. HMMs are used in several applications, including speech recognition [Rabiner 1989], natural language processing (NLP) [Manning & Schütze 1999], and the analysis of biological processes, such as protein topology [Krogh *et al.* 2001].

The predominant methods for estimating HMM parameters rely on local search heuristics. The *Baum-Welch (BW) algorithm* [Baum *et al.* 1970] and related methods are widely used. They utilize the *expectation-maximization (EM) algorithm* to find maximum likelihood parameter estimates ([Dempster *et al.* 1977] and [Bilmes 1998]). Unfortunately, these methods are prone to local optima and slow convergence [Redner & Walker 1984], and oftentimes suffer from high computational and sample complexity.

An alternative approach to the HMM learning problem is the *method of moments*. This parameter estimation technique was first proposed by [Pearson 1894]. The underlying idea is to sample empirical moments from experimental data, and then to find model parameters that yield expected values equal to the sampled quantities. Efficient learning algorithms for HMMs using this approach were enabled by a spectral decomposition technique introduced by [Chang 1996]. These algorithms can also be referred to as *spectral methods*. In contrast to the iterative approach used by the local search heuristics, spectral methods are not susceptible to local optima. Especially when learning parameter estimates from large observation spaces, this is an important advantage.

In this work at hand, the use of spectral methods for learning hidden Markov models is investigated. Spectral algorithms by [Hsu *et al.* 2012]

and [Anandkumar *et al.* 2012b] are introduced and characterized using synthetic experimental data.
Considering one of the most basic real-world settings, a measurement setup is likely to produce experimental data discretized in time with real-valued data. Assuming a hidden Markov model is to be used to describe the underlying process, it is important to find an appropriate way to prepare the experimental data for the use in the learning algorithms. Besides characterizing the spectral algorithms, this work will therefore also discuss how the algorithms could be used to obtain parameter estimates from sequences of real-valued observations.
For further reading on learning problems and machine learning in general, the book [Bishop 2006] provides a comprehensive treatment and can be recommended.

1.1 Outline

In this introductory chapter, an overview over the state of research is given, followed by an introduction to hidden Markov models.
The second chapter presents the two spectral learning algorithms by [Hsu *et al.* 2012] and [Anandkumar *et al.* 2012b]. For the latter algorithm, a detailed derivation will be given. The last section of the second chapter addresses how real-valued emissions can be used as an input for the learning algorithms.
The third chapter presents and discusses the results of the experiments conducted in this work. First, the two spectral algorithms will be compared with each other. In a second part, the algorithm by [Anandkumar *et al.* 2012b] will be compared with the Baum-Welch algorithm. In the last section, an experimental analysis of the use of real-valued emissions as the input data for spectral algorithms is given.
Chapter four provides the conclusion of this work and gives an outlook on possible future work.

1.2 State of research

Spectral methods for learning hidden Markov models were only established rather recently. Even though the underlying method of moments is a rather old concept, a major issue was that the estimation of high-order moments resulted in a high computational and sample complexity.

A spectral decomposition technique proposed by [Chang 1996] can be considered a milestone in enabling efficient spectral learning algorithms for HMMs. The method recovers model parameters using moments only up to third-order. The work also showed that second-order moments are not sufficient to fully identify the parameters of Markov models.
Based on this spectral decomposition, [Mossel & Roch 2005] proposed an efficient learning algorithm for phylogenic trees and HMMs. Compared to this algorithm by [Mossel & Roch 2005], the two algorithms that are under consideration in the work at hand share a similar rank condition on the emission and transition matrix of a HMM, and use similar empirical quantities computed from the observations for learning. [Mossel & Roch 2005] require the emission and transition matrix to have the same dimension.
The algorithm by [Hsu *et al.* 2012], which is the subject of this work, relaxed this assumption, allowing a larger observation space. The emission and transition matrix of the HMM are still assumed to have full rank. In favor of a higher sample-efficiency, this algorithm does not yield empirical estimates for the emission and transition matrices though. Instead, a so-called observable representation of probabilistic quantities of a HMM is given. The work by [Hsu *et al.* 2012] was initially published to the *arXiv.org* repository in 2009 and resulted in a work by [Siddiqi *et al.* 2009].
[Siddiqi *et al.* 2009] relaxed the assumption for the transition matrix to have full rank, resulting in an extended application of the algorithm by [Hsu *et al.* 2012] to the case of so-called *Reduced-Rank Hidden Markov Models* (RR-HMMs). In RR-HMMs, the dynamics are assumed to be evolving in a subspace smaller than the dimension of the transition matrix.
The second algorithm that is the subject of the work at hand was proposed by [Anandkumar *et al.* 2012b]. The algorithm follows the same approach as the algorithms mentioned above, and assumes the same full-rank condition. The learning algorithm for HMMs is proposed as a special case of a general learning algorithm for multi-view mixture models.
The works by [Mossel & Roch 2005], [Anandkumar *et al.* 2012b] and [Hsu *et al.* 2012] (plus the related algorithm by [Siddiqi *et al.* 2009]) utilize some variant of a diagonalization of a collection of similar matrices obtained from a tensor. The work by [Anandkumar *et al.* 2012a] addresses the issue of separation among eigenvalues when calculating model parameters from the diagonalization, which can lead to instability. [Anandkumar *et al.* 2012a] proposes a tensor power method allowing a more stable calculation of model parameters from an experimentally obtained tensor quantity. This theoretical work addresses latent variable models in general, listing hidden Markov models as a possible application. An explicit learning algorithm for HMMs is not formulated.

1.3 Notation

For two vectors $\vec{v}$ and $\vec{w}$, we denote the standard inner product by $\langle \vec{v}, \vec{w} \rangle = \vec{v}^\top \vec{w}$.
Let $\Delta^{n-1} := \{(p_1, p_2, \ldots, p_n) \in \mathbb{R}^n \,|\, p_i \geq 0 \,\forall\, i, \sum_{i=1}^n p_i = 1\}$ denote the probability simplex in $\mathbb{R}^n$.
For a matrix $A \in \mathbb{R}^{m \times n}$, we let $\|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2}$ denote its Frobenius norm and A^+ its Moore–Penrose pseudoinverse.
Let $\vec{u} \in \mathbb{R}^n$ and $B \in \mathbb{R}^{m \times n}$ be the analytical quantities of a vector and a matrix. We then denote their empirical estimates from experimental data by $\hat{\vec{u}} \in \mathbb{R}^n$ and $\hat{B} \in \mathbb{R}^{m \times n}$, respectively.

1.4 Hidden Markov models

Hidden Markov models were developed by Leonard E. Baum and coworkers in [Baum & Petrie 1966] and subsequent works. In this section, the mathematical foundations are introduced. In a first step, the concept of a Markov process is presented. The introduction of latent variables then leads to the hidden Markov model. Model parameters are introduced as probabilistic measures, leading to the formulation of the learning problem. The mathematical presentation in this section follows [Bishop 2006].

1.4.1 Markov process

Let $p(\cdot)$ denote a probability distribution function. Let there be a sequence of L observations, with s_t denoting the observation at time $t \in \{1, ..., L\}$. The observation can be a vector or scalar. The joint distribution of a sequence of L observations is given by

$$p(s_1, \ldots, s_L) = p(s_1) \prod_{t=2}^{L} p(s_t | s_1, \ldots, s_{t-1}). \tag{1.1}$$

Assuming the Markov property, the current observation in a sequence only depends on the previous one. This is also known as a *Markov chain*. For a sequence of observations, the joint distribution can therefore be written as

$$p(s_1, \ldots, s_L) = p(s_1) \prod_{t=2}^{L} p(s_t | s_{t-1}). \tag{1.2}$$

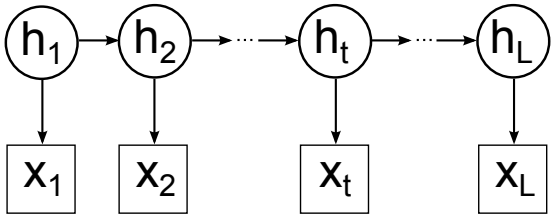

Figure 1.1: A graphical structure representing a hidden Markov model: a sequence of discrete hidden states $\{h_t\}_{t=1,\dots,L}$ with the corresponding emissions $\{x_t\}_{t=1,\dots,L}$.

One can verify that, given previous observations up to time t, this yields the following conditional distribution for an observation s_t:

$$p(s_t|s_1,\dots,s_{t-1}) = p(s_t|s_{t-1}). \tag{1.3}$$

In a hidden Markov model, we now assume an underlying Markov process that cannot be observed. While the states are hidden, they emit an output which is observable. In the model, each state has a corresponding probability distribution over the possible emissions. The hidden states are interpreted as discrete *latent variables* which are not directly observable. These latent variables are then assumed to form a Markov chain. The possible emissions can be either discrete or continuous.

1.4.2 Hidden states as latent variables

We assume that the underlying states form an unobservable Markov chain. We denote a hidden state at time t by h_t. The possible discrete or continuous emissions by this state are denoted by x_t. The resulting graphical structure is depicted in Figure 1.1.
This setup now allows the parameterization of a hidden Markov model only by the transition probability between states, the probability distribution over possible emissions by each state, and the initial distribution of states.

The joint distribution for a hidden Markov model is given by

$$p(x_1,...,x_L,h_1,...,h_L) = p(h_1)\left[\prod_{t=2}^{L} p(h_{t+1}|h_t)\right]\prod_{t=1}^{L} p(x_t|h_t). \tag{1.4}$$

1.4.3 State transitions

Let there be k possible hidden states. Since the latent variables in a HMM are discrete, we will denote the set of hidden states by $[k] = \{1, ..., k\}$. In a HMM, the transition probabilities between two states do not evolve with time. Therefore we can introduce a *state transition probability matrix* . Let $T \in \mathbb{R}^{k \times k}$ be such a matrix with

$$T_{ij} = \Pr\left[h_{t+1} = i | h_t = j\right], \quad i, j \in [k]\,. \tag{1.5}$$

T_{ij} is the probability for the latent variable h_{t+1} to be in the state $i \in [k]$ at time $t+1$ after having been in state $j \in [k]$ at time t. The elements of the matrix are probabilities, and therefore $T_{ij} \in [0, 1]$ with $\sum_{i=1}^{k} T_{ij} = 1,\ \forall j \in [k]$. This results in $k(k-1)$ independent parameters for the transition matrix.

For the initial state h_1, the *initial state distribution* is denoted by the vector $\vec{\pi} \in \mathbb{R}^k$, with

$$\pi_j = \Pr\left[h_1 = j\right]. \tag{1.6}$$

1.4.4 Emissions

The emissions of a hidden Markov model can be parameterized by an *emission probability matrix*. The simplest case is for *categorical* emissions. Here the emission probability matrix can be intuitively understood as a probability table.
Let there be d discrete observable emissions, with $d \geq k$. The set of discrete emissions can be written as $[d] = \{1, ..., d\}$. The emission $x_t = i, i \in [d]$ was emitted by the hidden state $h_t = j, j \in [k]$ with the probability $\Pr\left[x_t = i | h_t = j\right]$. This now allows the introduction of the emission probability matrix $O \in \mathbb{R}^{d \times k}$, with

$$O_{ij} = \Pr\left[x_t = i | h_t = j\right], \quad i \in [d]\,, j \in [k]\,. \tag{1.7}$$

Another case are *vector-valued* emissions. Here an emission at time t is represented by a random vector $\vec{x}_t$ with values in $\mathbb{R}^d$. In this case we can introduce the emission matrix as

$$O = [\vec{o}_1 | \vec{o}_2 | \ldots | \vec{o}_k] \in \mathbb{R}^{d \times k}.$$

The columns of the matrix denote the conditional means of the emissions $\vec{x}_t$ at time t, given the corresponding discrete hidden state h_t:

$$\mathbb{E}\left[\vec{x}_t | h_t = j\right] = O\vec{e}_j = \vec{o}_j, \quad j \in [k]\,,$$

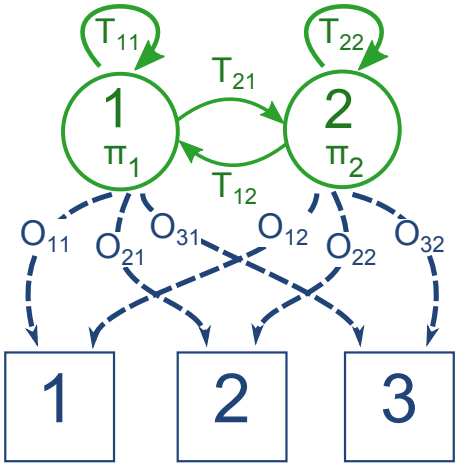

Figure 1.2: A two-state HMM (states 1 and 2) with three possible emissions (1, 2 and 3). The initial state distribution is indicated inside the circles of the corresponding states (π_1 and π_2). Solid arrows indicate transitions between hidden states, dashed arrows indicate emissions of the possible observations.

where $\vec{e}_j$ is the j-th vector of the d-dimensional standard basis. The conditional mean of an emission could for example represent the mean of a multivariate Gaussian.
The words *emission* and *observation* are often used interchangeably. In this work, we understand that in the generative process, a hidden state emits an output which can be observed. This output can be interpreted both as the emission of the hidden state, or the observation of this emission. Especially in the experimental context, it can be considered more appropriate to talk about an observation, rather than an emission.
Both terms will be used in this work. The term *observation* will be preferred in the context of experimental data, while the term *emission* will be used when referring to the generative process.

1.4.5 Example

As an example, we now consider a two-state hidden Markov model as depicted in Figure 1.2. The underlying system can be in the hidden states 1 or 2. At a given timestep, the system then emits one of the three observations 1,2 or 3, depending on which hidden state it is in. Using equations (1.5), (1.6) and (1.7), the model can then be parameterized in the following way: let $T \in \mathbb{R}^{2\times2}$, $\vec{\pi} \in \mathbb{R}^2$ and $O \in \mathbb{R}^{3\times2}$, then

$$T = \begin{pmatrix} T_{11} & T_{12} \\ T_{21} & T_{22} \end{pmatrix} = \begin{pmatrix} 9/10 & 3/10 \\ 1/10 & 7/10 \end{pmatrix}, \tag{1.8}$$

$$\vec{\pi} = \begin{pmatrix} \pi_1 \\ \pi_2 \end{pmatrix} = \begin{pmatrix} 4/5 \\ 1/5 \end{pmatrix} \tag{1.9}$$

and

$$O = \begin{pmatrix} O_{11} & O_{12} \\ O_{21} & O_{22} \\ O_{31} & O_{32} \end{pmatrix} = \begin{pmatrix} 1/4 & 8/10 \\ 1/2 & 1/10 \\ 1/4 & 1/10 \end{pmatrix}. \tag{1.10}$$

The numerical values given are exemplary values. A HMM parameterized by these matrices will be one of the models used in our analysis of the learning algorithms in Chapter 3.

1.4.6 Learning problem

The learning problem for hidden Markov models is generally an instance of *unsupervised learning*. In unsupervised learning, the experimental data is unlabeled, i.e. the desired output is not known. [1]
A HMM can be characterized by two model parameters, the number of states k and the number (or dimension) of emissions d, and three probability measures, T, O and $\vec{\pi}$. [2] In a more compact notation, the probability measures are often grouped as a complete parameter set of the model:

$$\lambda = (T, O, \vec{\pi}) .$$

Let $p(x_1, \ldots, x_L|\lambda)$ be the probability of observing the sequence $(x_1, \ldots, x_L)$, provided the model parameters λ. For a hidden Markov model to be useful in a real-world application, [Rabiner 1989] proposed the following three basic problems of interest:

1. Given a sequence of observations $(x_1, \ldots, x_L)$ and a model $\lambda = (T, O, \vec{\pi})$, how can $p(x_1, \ldots, x_L|\lambda)$ be computed efficiently?

2. Given a sequence of observations $(x_1, \ldots, x_L)$ and the model λ, how do we choose a (scientifically) meaningful or optimal corresponding sequence of states $(h_1, \ldots, h_L)$?

3. How can the model parameters $\lambda = (T, O, \vec{\pi})$ be chosen such that they best account for the observed signal (i.e. maximize $p(x_1, \ldots, x_L|\lambda)$)?

[1] For further information, refer to [Bishop 2006].
[2] We only consider the case of a homogeneous model where the parameters T,O,$\vec{\pi}$ represent all corresponding conditional probabilities for the respective states and emissions.

This work will address the last of the three basic problems. In our setting, we will assume we are given a series of observations. These observations will then be used to learn the model parameters λ of the hidden Markov model.

Chapter 2

Theory

This chapter will introduce two spectral algorithms. The algorithm by [Hsu *et al.* 2012] will be presented in Section 2.1. The presentation will be kept rather brief, since a comprehensive derivation is given in the original work. Furthermore, we will refrain from introducing the concept of observable representations associated with this algorithm, since it will not be relevant for our considerations.
The algorithm by [Anandkumar *et al.* 2012b] will be presented in Section 2.2. The original work does not explicitly formulate a learning algorithm for hidden Markov models. Therefore, we will give a comprehensive derivation of the mathematical considerations leading to the formulation of the spectral learning algorithm for HMMs.
Section 2.3 will discuss real-valued emissions.

2.1 Spectral algorithm by Hsu et al.

Here, we will introduce the spectral learning algorithm proposed by [Hsu *et al.* 2012] for HMMs with discrete emissions. The algorithm in its basic form does not explicitly learn parameter estimates for the transition and emission probability matrices. It only learns a so-called observable representation allowing the estimation of the probability of observing a given sequence, and the approximation of the conditional probability of a future observation.
[Hsu *et al.* 2012] points out a method to extend the algorithm to allow an estimation of the state transition probability matrix and emission probability matrix. Therefore, we will introduce the algorithm disregarding its observable representation functionality, and include the method to recover estimates for the state transition probability and emission probability ma-

trices. It has to be noted, however, that the authors in [Hsu *et al.* 2012] explicitly state that this additional estimation technique is to be considered generally unstable.
Let our HMM be parameterized as in Section 1.4, with our observations being discrete. We then have k hidden states denoted by $[k] = \{1, \ldots, k\}$ and d discrete observations denoted by $[d] = \{1, \ldots, d\}$ with $d \geq k$. Our state transition probability matrix $T \in \mathbb{R}^{k \times k}$ is then given by (1.5), the emission probability matrix $O \in \mathbb{R}^{d \times k}$ by (1.7) and the initial state distribution $\vec{\pi} \in \mathbb{R}^k$ by (1.6).
We assume the following condition on the HMM:

Condition 2.1.1. (HMM rank condition). $\pi_j > 0$ for all $j \in [k]$, and $\operatorname{rank}(O) = \operatorname{rank}(T) = k$.

Furthermore, we define the following vector and matrix quantities:

Definition 2.1.1. Let $P_1 \in \mathbb{R}^d$, $P_{2,1} \in \mathbb{R}^{d \times d}$, $P_{3,1} \in \mathbb{R}^{d \times d}$ and $P_{3,r,1} \in \mathbb{R}^{d \times d}$ with

$$
\begin{aligned}
&[P_1]_i = \Pr[x_1 = i], \\
&[P_{2,1}]_{ij} = \Pr[x_2 = i, x_1 = j], \\
&[P_{3,1}]_{ij} = \Pr[x_3 = i, x_1 = j], \text{and} \\
&[P_{3,r,1}]_{ij} = \Pr[x_3 = i, x_2 = r, x_1 = j] \quad \forall r \in [d].
\end{aligned}
$$

These quantities are marginal probabilities of singletons, pairs and triples of observations. These quantities will later be sampled from the experimental data.
The algorithm utilizes a singular value decomposition of the empirically sampled estimation matrix of $P_{2,1}$. Let the matrix of the top k left singular vectors of $P_{2,1}$ be denoted by $U \in \mathbb{R}^{d \times k}$. We assume that the matrix obeys the following condition:

Condition 2.1.2. (Invertibility condition). $U^\top O$ is invertible.

2.1.1 Algorithm learnHKZ

The algorithm stated in [Hsu *et al.* 2012] does not explicitly learn the state transition probability matrix and emission probability matrix of a HMM. However, the work also proposes a technique to learn these parameters.

The resulting spectral learning algorithm is given by

Algorithm 1: LEARNHKZ [Hsu *et al.* 2012]

Data: N triples of observations, k - number of states , d - number of observations

Result: Hidden Markov model parameterized by $\hat{O}$, $\hat{T}$ and $\hat{\vec{\pi}}$

1. Form the empirical estimates $\hat{P}_1 \in \mathbb{R}^d$, $\hat{P}_{2,1} \in \mathbb{R}^{d\times d}$, $\hat{P}_{3,1} \in \mathbb{R}^{d\times d}$ and $\hat{P}_{3,r,1} \in \mathbb{R}^{d\times d}\ \forall r \in [d]$ using the N triples of observations.

2. For the k largest singular values of $\hat{P}_{2,1}$, find the matrix $\hat{U} \in \mathbb{R}^{d\times k}$ of left singular vectors.

3. For all $r \in [d]$, find the eigenvalues $\lambda_{r,1}, \ldots, \lambda_{r,k}$ of

$$\left(\hat{U}^\top \hat{P}_{3,r,1}\right)\left(\hat{U}^\top \hat{P}_{3,1}\right)^+$$

4. Return the empirical estimate of the matrix $\hat{O}$ given by:

$$\hat{O} = \begin{pmatrix} \lambda_{1,1} & \lambda_{1,2} & \cdots & \lambda_{1,k} \\ \lambda_{2,1} & \lambda_{2,2} & \cdots & \lambda_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{d,1} & \lambda_{d,2} & \cdots & \lambda_{d,k} \end{pmatrix}$$

5. Return $\hat{\vec{\pi}} = \hat{O}^+ \hat{P}_1$.

6. Return $\hat{T} = \hat{O}^+ \hat{P}_{2,1} (\hat{O}^+)^\top \mathrm{diag}(\hat{\vec{\pi}})^{-1}$.

The underlying basic algorithm LEARNHMM by [Hsu *et al.* 2012] has polynomial sample and computational complexity. Furthermore, assuming Conditions 2.1.1 and 2.1.2, the basic algorithm yields a provably correct result. However, the algorithm does not ensure that the predicted probabilities lie in the range $[0, 1]$ [Hsu *et al.* 2012].

2.2 Spectral algorithm by Anandkumar et al.

This section will establish a learning algorithm as proposed in [Anandkumar *et al.* 2012b], including important mathematical derivations.
A simple bag-of-words model will be used to introduce cross moments and

their representation using a tensor notation. This will then lead to the formulation of multi-view models and their application for hidden Markov models. Since the work by [Anandkumar *et al.* 2012b] does not explicitly state a learning algorithm for HMMs, part 2.2.3 will show important mathematical derivations. These derivations then lead to the spectral learning algorithm presented in part 2.2.4.
The presentation closely follows [Anandkumar *et al.* 2012b] and [Anandkumar *et al.* 2012a]. The proofs can be found in Appendix A.

2.2.1 Bag-of-words

The *bag-of-words model* is a simple model where a document is assumed to be composed of a set of words in exchangeable order. Word ordering and grammar are therefore disregarded, but multiplicity of words is observed. Intuitively, in this model a document can be imagined as a bag full of unordered words.
The generative process can be described in two steps. First, the topic for a document is chosen as specified by a multinomial distribution out of the discrete set of distinct topics. Then, for the given topic, the document's words are drawn independently from the set of words, governed by another multinomial distribution.
Let there be a body of k distinct topics a document can have. Let d be the number of distinct words in the vocabulary and let l be the number of words in a document. We assume $l \geq 3$.
The latent variable h is identified with the topic of a given document. The probability for a document to have the topic $j \in [k]$ is given by

$$\Pr[h = j] = w_j, \quad j \in [k],$$

where $\vec{w} := (w_1, w_2, ..., w_k) \in \Delta^{k-1}$ is a probability vector governing the distribution of topics.
For a given topic h, the probability vector $\vec{\mu}_h \in \Delta^{k-1}$ specifies a multinomial distribution from which the l words of the documents are drawn. A practical way to represent the words is by using d-dimensional vectors.
Let the set of l words in a document be $\vec{x}_1, \vec{x}_2, ..., \vec{x}_l \in \mathbb{R}^d$. We then set

$$\vec{x}_t = \vec{e}_i \Leftrightarrow \text{the } t\text{-th word in the document is } i, \quad t \in [l],$$

where $\vec{e}_1, \vec{e}_2, ..., \vec{e}_d$ is the standard coordinate basis for $\mathbb{R}^d$.
This vector representation can now be used to intuitively describe joint probabilities over words.

Let $\mathbb{E}[\vec{x}_1 \otimes \vec{x}_2]$ be the matrix of joint probabilities over two words. The matrix is then given by

$$\begin{aligned}\mathbb{E}[\vec{x}_1 \otimes \vec{x}_2] &= \sum_{1 \leq i,j \leq d} \Pr[\vec{x}_1 = \vec{e}_i, \vec{x}_2 = \vec{e}_j]\, \vec{e}_i \otimes \vec{e}_j \\ &= \sum_{1 \leq i,j \leq d} \Pr\left[1^{\text{st}} \text{ word} = i, 2^{\text{nd}} \text{ word} = j\right] \vec{e}_i \otimes \vec{e}_j.\end{aligned}$$

Therefore the (i,j)-th entry of the matrix $\mathbb{E}[\vec{x}_1 \otimes \vec{x}_2]$ of joint probabilities is $\Pr\left[1^{\text{st}} \text{ word} = i, 2^{\text{nd}} \text{ word} = j\right]$. This leads to a general tensor formulation for the joint probabilities over l words:

$$\mathbb{E}[\vec{x}_1 \otimes \vec{x}_2 \otimes \cdots \otimes \vec{x}_l] =$$
$$\sum_{1 \leq i_1, i_2, \ldots, i_l \leq d} \Pr\left[1^{\text{st}} \text{ word} = i_1, 2^{\text{nd}} \text{ word} = i_2, \ldots, l^{\text{th}} \text{ word} = i_l\right] \vec{e}_{i_1} \otimes \vec{e}_{i_2} \otimes \cdots \otimes \vec{e}_{i_l}.$$

The $(i_1, i_2, ..., i_l)$-th entry of the tensor $\mathbb{E}[\vec{x}_1 \otimes \vec{x}_2 \otimes \cdots \otimes \vec{x}_l]$ corresponds to $\Pr\left[1^{\text{st}} \text{ word} = i_1, 2^{\text{nd}} \text{ word} = i_2, ..., l^{\text{th}} \text{ word} = i_l\right]$.
An important result of this consideration is that the estimation of cross moments can be conducted by the estimation of joint probabilities of a sequence of words. As an example, the cross moment for $\vec{x}_1 \otimes \vec{x}_2 \otimes \vec{x}_3$ can be estimated by sampling the joint probabilities of the first three words over all documents.
Let us now consider the conditional expectation of $\vec{x}_t$ given a topic $h = j$, $j \in [k]$:

$$\begin{aligned}\mathbb{E}[\vec{x}_t | h = j] &= \sum_{i=1}^{d} \Pr[t\text{-th word} = i | h = j]\, \vec{e}_i \\ &= \sum_{i=1}^{d} [\vec{\mu}_j]_i\, \vec{e}_i \\ &= \vec{\mu}_j.\end{aligned}$$

The conditional expectation of a word $\vec{x}_t$ given a document topic $h = j$ is therefore equal the probability vector $\vec{\mu}_j \in \Delta^{k-1}$, $j \in [k]$. This is another important property that will be used in later considerations.

2.2.2 Multi-view mixture models

Coming from the bag-of-words model introduced in the previous section 2.2.1, a generalization now leads to the concept of *multi-view models*. Similarly to the bag-of-words model, our observed variables $\vec{x}_1, \vec{x}_2, ..., \vec{x}_l$ are

assumed to be conditionally independent given a latent variable h. However, now the conditional distributions of the observations $\vec{x}_t$, $t \in [l]$ are not required to be identical anymore.
In a bag-of-words model, the topic of a document does not change while the words are picked according to their multinomial distribution. When considering the observations of a hidden Markov model however, the conditional distribution of an observation depends on the underlying hidden state. Since a HMM passes through a hidden state sequence, the state of the system might have changed from one observation to the subsequent one.

General setting

We now introduce the general setting for a multi-view mixture model. This setting needs to be established in order to properly introduce a learning algorithm as proposed by [Anandkumar *et al.* 2012b]. In terms of its application to HMMs, the concept might not seem to be very intuitive at first. In the course of the formulation of the learning algorithm for HMMs in 2.2.3, the relationship between HMMs and the general quantities introduced here should become clear. To allow for a better understanding of the application of the multi-view mixture model to HMMs, the presentation here will be slightly different to [Anandkumar *et al.* 2012b].
We restrict our discussion to x_t with $t \in \{1, 2, 3\}$. This results in an instance of a three-view mixture model which will be the starting point for our considerations regarding hidden Markov models.

Definition 2.2.1. Let k be the number of mixture components. [1] Let l be the number of views, assuming $l \geq 3$. Let $\vec{w} = (w_1, w_2, ..., w_k) \in \Delta^{k-1}$ be a vector of mixing weights, and let $h \in [k]$ be a hidden discrete random variable with $\Pr[h = j] = w_j$, $\forall j \in [k]$. Let $\vec{x}_1, \vec{x}_2, \vec{x}_3 \in \mathbb{R}^d$ be random vectors that are conditionally independent given h. We define the conditional means of these vectors as

$$\vec{\mu}_{t,j} := \mathbb{E}[\vec{x}_t | h = j], \quad j \in [k], t \in \{1, 2, 3\}.$$

and $M_t \in \mathbb{R}^{d \times k}, t \in \{1, 2, 3\}$ as the matrix of conditional means, with the j-th column being $\vec{\mu}_{t,j}$:

$$M_t := [\vec{\mu}_{t,1} | \vec{\mu}_{t,2} | \dots | \vec{\mu}_{t,k}].$$

[1] In the case of HMMs, this will later correspond to the number of states.

2.2.3 Derivations

The algorithm proposed in [Anandkumar *et al.* 2012b] is formulated as a general method for learning multi-view mixture models. This section elaborates important considerations for the formulation as a spectral algorithm for HMMs. The resulting learning algorithm will then be presented in 2.2.4. Using the setting introduced in Definition 2.2.1 we assume the following condition.

Condition 2.2.1. Non-degeneracy.
$w_j > 0$ for all $j \in [k]$, and $\operatorname{rank}(M_t) = k$ for all $t \in \{1, 2, 3\}$.

To make the transition from the general setting to hidden Markov models, we now relate the quantities $M_t, t \in \{1, 2, 3\}$ to the parameters of a hidden Markov model.

Proposition 1. *[Anandkumar et al. 2012b] If the hidden variable h from the multi-view mixture model (Section 2.2.2) is identified with the second hidden state h_2, then $\{\vec{x}_1, \vec{x}_2, \vec{x}_3\}$ are conditionally independent given h, and the parameters of the resulting three-view mixture model on $(h, \vec{x}_1, \vec{x}_2, \vec{x}_3)$ are*

$$\vec{w} = T\vec{\pi}, \quad M_1 = O \operatorname{diag}(\vec{w})\, T^\top \operatorname{diag}(T\vec{\pi})^{-1}, \quad M_2 = O, \quad M_3 = OT.$$

For the following derivations we will however stick to the general notation using M_1, M_2 and M_3 at first, as the derivations are also valid for multi-view models in general. Furthermore this simplifies and shortens notation.

Definition 2.2.2. Second- and third-order moments.
We define the matrix of second-order moments $P_{3,1} \in \mathbb{R}^{d \times d}$ and the tensor of third-order moments $P_{3,1,2} \in \mathbb{R}^{d \times d \times d}$ by

$$P_{3,1} := \mathbb{E}\left[\vec{x}_3 \otimes \vec{x}_1\right] \quad \text{and} \quad P_{3,1,2} := \mathbb{E}\left[\vec{x}_3 \otimes \vec{x}_1 \otimes \vec{x}_2\right]. \tag{2.1}$$

The tensor $P_{3,1,2}$ can also be regarded as a linear operator $P_{3,1,2} : \mathbb{R}^d \to \mathbb{R}^{d \times d}$ given by

$$P_{3,1,2}(\vec{\eta}) := \mathbb{E}\left[(\vec{x}_3 \otimes \vec{x}_1)\, \langle \vec{\eta}, \vec{x}_2 \rangle\right], \tag{2.2}$$

where $\langle \vec{x}, \vec{y} \rangle = \vec{x}^\top \vec{y}$ denotes the standard inner product between two vectors.
For a vector $\vec{\eta} = (\eta_1, ..., \eta_d)$ the (i, j)-th entry of the matrix $P_{3,1,2}(\vec{\eta})$ is then given by

$$P_{3,1,2}(\vec{\eta})_{i,j} = \sum_{\alpha=1}^{d} \eta_\alpha \left[P_{3,1,2}\right]_{i,j,\alpha}, \tag{2.3}$$

where $[P_{3,1,2}]_{i,j,\alpha}$ is the (i, j, α)-th element of the third-order moment tensor.

The second- and third-order moments in Definition 2.2.2 can be related to the model parameters for multi-view mixture models. These relations will play an important role in formulating the parameter estimation algorithm.
The following lemma relates the moments defined above to the model parameters $M_t, t \in \{1, 2, 3\}$.

Lemma 1. *Let $P_{3,1}$ and $P_{3,1,2}$ be defined as in Definition 2.2.2. Then*

$$P_{3,1} = M_3 \operatorname{diag}(\vec{w})\, M_1^\top \quad \text{and} \quad P_{3,1,2} = M_3 \operatorname{diag}(M_2^\top \vec{\eta}) \operatorname{diag}(\vec{w})\, M_1^\top \quad \forall \vec{\eta} \in \mathbb{R}^d.$$

Given a sequence of observations, it is possible to efficiently sample the pair-wise and triple-wise probabilities to form empirical estimates for the second- and third-order moments. Lemma 1 now implies that these moments can be used to reveal the underlying model parameters. This is done using a so-called *observable operator*. Observable operators were introduced by [Jaeger 2000] and are also used in the HMM learning algorithm by [Hsu *et al.* 2012].
The following important lemma introduces such an observable operator, a quantity that is only a function of joint probabilities of observable variables. Furthermore, the lemma relates this operator to our model parameters.

Lemma 2. *[Anandkumar et al. 2012b] Assume Condition 2.2.1. For $t \in \{1, 2, 3\}$, let U_t be a matrix such that $U_t^\top M_t$ is invertible. Then $U_3^\top P_{3,1} U_1$ is invertible, and for all $\vec{\eta} \in \mathbb{R}^d$, the observable operator $B_{3,1,2} \in \mathbb{R}^{k \times k}$, given by*

$$B_{3,1,2}(\vec{\eta}) := (U_3^\top P_{3,1,2}(\vec{\eta}) U_1)(U_3^\top P_{3,1} U_1)^{-1} \tag{2.4}$$

satisfies

$$B_{3,1,2}(\vec{\eta}) = \left(U_3^\top M_3\right) \operatorname{diag}\left(M_2^\top \vec{\eta}\right) \left(U_3^\top M_3\right)^{-1}. \tag{2.5}$$

In particular, the k roots of the polynomial $\lambda \mapsto \det\left(B_{3,1,2}(\vec{\eta}) - \lambda I\right)$ are $\{\langle \vec{\eta}, \vec{\mu}_{2,j} \rangle : j \in [k]\}$.

The matrix $B_{3,1,2}(\vec{\eta})$ is therefore similar to the diagonal matrix $\operatorname{diag}\left(M_2^\top \vec{\eta}\right)$, and the columns of $U_3^\top M_3$ are the eigenvectors of $B_{3,1,2}(\vec{\eta})$. Considering the case where $\vec{\eta} = \vec{e}_i, i \in [d]$, this would now relate our desired model parameters to our observable operator matrix. In absence of prior knowledge, these vectors $\vec{\eta} = \vec{e}_i$ are not known though. It is, however, possible to obtain the model parameters by applying different observable operators to different vectors $\vec{\eta}$:

Lemma 3. *[Anandkumar et al. 2012b] Consider the setting and definitions from Lemma 2. Let $\Theta \in \mathbb{R}^{k\times k}$ be an invertible matrix, and let $\vec{\theta}_i \in \mathbb{R}^k$ be its i-th row. For all $i \in [k]$, let $\lambda_{i,1}, \lambda_{i,2}, ..., \lambda_{i,k}$ denote the k eigenvalues of $B_{3,1,2}(U_2\vec{\theta}_i)$ in the order specified by the matrix of right eigenvectors $U_3^\top M_3$. Let $L \in \mathbb{R}^{k\times k}$ be the matrix whose (i,j)-th entry is $\lambda_{i,j}$. Then*

$$\Theta U_2^\top M_2 = L. \tag{2.6}$$

This lemma now shows that the unknown parameters M_2 can be obtained by the solution of the linear system (2.6). The elements of the matrix L are the roots of the k-th degree polynomial of the observable operator matrices derived from the empirically observed second- and third-order moments. Proposition 1 now states that $M_2 = O$, so we now have established a method to learn the emission probability matrix:

$$O = U_2\Theta^{-1}L. \tag{2.7}$$

The obtained matrix O together with Lemma 2 and Proposition 1 now enables us to calculate an estimate for the state transition probability matrix T. Using Equation (2.5) and inserting the quantities M_3 and M_2 from Proposition 1 we obtain

$$\begin{aligned} B_{3,1,2}(\vec{\eta}) &= \left(U_3^\top M_3\right) \operatorname{diag}\left(M_2^\top\vec{\eta}\right) \left(U_3^\top M_3\right)^{-1} \\ &= \underbrace{\left(U_3^\top OT\right)}_{=:R} \operatorname{diag}\left(O^\top\vec{\eta}\right) \left(U_3^\top OT\right)^{-1} \\ &= R \operatorname{diag}\left(O^\top\vec{\eta}\right) R^{-1}, \end{aligned}$$

where R is the matrix of right eigenvectors of $B_{3,1,2}(\vec{\eta})$. Therefore, our state transition probability matrix is given by

$$T = \left(U_3^\top O\right)^{-1} R. \tag{2.8}$$

2.2.4 Algorithm learnAHK

The spectral learning algorithm is now the result of the relations revealed by the derivations in 2.2.3. The lemmas suggest the following estimation procedure:
Proposition 1 shows how the parameters of a hidden Markov model relate to the moments in a multi-view (more specifically, three-view) mixture model. Lemma 1 relates these moments to quantities that can be obtained empirically, given a sequence of observations. The observable operators introduced in Lemma 2 can then be used to reveal our model parameters by

a simultaneous diagonalization as implied by Lemma 3.
Note that due to the nature of the singular value decomposition, the ordering of the columns of the estimated quantities is subject to permutation. The estimated transition probability matrix is only calculated up to a scaling of the columns. The implementation of the algorithm in this work performed a basic normalization to obtain transition probabilities that sum to 1.
The framework of the algorithm is suitable for both categorical observations and continuous multivariate observations.

Algorithm 2: LEARNAHK [Anandkumar *et al.* 2012b]

Data: N triples of observations, k - number of states , d - number of observations

Result: Hidden Markov model parameterized by $\hat{O}$ and $\hat{T}$

1. Obtain empirical estimates for $\hat{P}_{3,1} \in \mathbb{R}^{k\times k}$, $\hat{P}_{3,2} \in \mathbb{R}^{k\times k}$ and $\hat{P}_{3,1,2} \in \mathbb{R}^{k\times k\times k}$ from the N observation triples.

2. For the largest k singular values of $\hat{P}_{3,1}$, find the matrix $\hat{U}_3 \in \mathbb{R}^{d\times k}$ of orthonormal left and the matrix $\hat{U}_1 \in \mathbb{R}^{d\times k}$ of orthonormal right singular vectors.
   For the largest k singular values of $\hat{P}_{3,2}$, find the matrix $\hat{U}_2 \in \mathbb{R}^{d\times k}$ of orthonormal right singular vectors.

3. Pick an invertible matrix $\Theta \in \mathbb{R}^{k\times k}$, with the i-th row denoted as $\vec{\theta}_i^\top \in \mathbb{R}^k$. In the absence of prior knowledge about $\hat{O}$, choose Θ to be a random rotation matrix.
   For all $i \in [k]$ calculate the matrix $\hat{P}_{3,1,2}(\hat{U}_3\vec{\theta}_i)$, with entries

$$\hat{P}_{3,1,2}(\hat{U}_3\vec{\theta}_i)_{k,l} = \sum_{\alpha=1}^{d} \left[\hat{U}_3\vec{\theta}_i\right]_\alpha [P_{3,1,2}]_{k,l,\alpha}\,.$$

   For all $i \in [k]$, form the matrices

$$\hat{B}_{3,1,2}(\hat{U}_3\vec{\theta}_i) = (\hat{U}_3^\top \hat{P}_{3,1,2}(\hat{U}_3\vec{\theta}_i)\hat{U}_1)(\hat{U}_3^\top \hat{P}_{3,1}\hat{U}_1)^{-1}$$

   Diagonalize the matrix $\hat{B}_{3,1,2}(\hat{U}_3^\top\vec{\theta}_1)$ such that

$$\hat{R}_3^{-1}\hat{B}_{3,1,2}(\hat{U}_3^\top\vec{\theta}_1)\hat{R}_3 = \mathrm{diag}\left(\hat{\lambda}_{1,1}, \hat{\lambda}_{1,2}, \ldots, \hat{\lambda}_{1,k}\right).$$

   (Fail if not possible.)

4. For each $i \in \{2, \ldots, k\}$, obtain diagonal entries $\hat{\lambda}_{i,1}, \hat{\lambda}_{i,2}, \ldots, \hat{\lambda}_{i,k}$ of $\hat{R}_3^{-1}\hat{B}_{3,1,2}(\hat{U}_3^\top\vec{\theta}_i)\hat{R}_3$ and form the matrix $\hat{L} \in \mathbb{R}^{k\times k}$ with the (i,j)-th entry being $\hat{\lambda}_{i,j}$.

5. Return $\hat{O} = \hat{U}_2\Theta^{-1}\hat{L}$.

6. Return $\hat{T} = \left(\hat{U}_3\hat{O}\right)^{-1}\hat{R}_3$, scale columns of $\hat{T}$.

2.3 Real-valued emissions

When considering experiments in real-world applications, measurement data is not likely to be categorical. Furthermore, measurements will suffer from both intrinsic noise in the underlying processes and extrinsic noise induced by the measurement setup.
Let there be a time series of real-valued emissions, $\{y_t\}_{t=1,2,\dots,L}$ with $y_t \in \mathbb{R}, \forall t$. Assume that the generative process of our measurement data can be described by a hidden Markov model. We now want to use the spectral algorithms (Section 2.1 and Section 2.2) to learn parameter estimates from this data. The algorithms do not allow the input of the real-valued sequences directly. Thus, the challenge is to find a representation for our real-valued emissions so that they can be used as an input for the spectral algorithms.
One possible representation would be to obtain a sequence of categorical emissions by binning the real-valued emissions. The approach evaluated here will be to bin by quantiles. Taking all the emissions available, quantiles can be calculated, leading to natural bin bounds.
A simple binning approach is introduced in this section. The approach will then be evaluated in Section 3.2.

2.3.1 Simple binning approach

Assume we are given a sequence of real-valued observations. Furthermore, assume that the dimension, or number, of emissions d of the hidden Markov model is given. One of the simplest approaches to prepare the data for the learning algorithms is to bin the observations to generate a sequence of categorical emissions. The major challenge is to find appropriate bin bounds. The approach pursued in this method is to use all available observations and calculate quantiles. These quantiles then yield natural bounds that are used for binning. The procedure is formulated in the following algorithm:

Algorithm 3: SIMPLEBINNING

Data: Sequence of real-valued observations, d - dimension of observations

Result: Sequence of binned observations

1. Using the real-valued observations, calculate the quantiles for $d-1$ evenly spaced cumulative probabilities $\left(\frac{1}{d}, \frac{2}{d}, \ldots, \frac{d-1}{d}\right)$
2. Observing the ordering of the sequence, bin each observation according to the bounds specified by the quantiles.
3. Return sequence of binned observations.

The sequence of categorical variables obtained by SIMPLEBINNING can then be used as the input for the spectral learning algorithms LEARNHKZ and LEARNAHK. This procedure allows the learning of HMM model parameters given a sequence of real-valued emissions.
There are, however, intrinsic problems arising from this simple binning approach. Quantiles are calculated using all available observations. The resulting bin bounds will divide the total of all real-valued observations into bins of equal size. In our categorical sequence resulting from the binning, each possible observation will therefore occur equally often across all elements of the sequence. When considering the HMM, this would correspond to the assumption that the cumulative probability of each of the possible emissions is the same. Without loss of generality, this assumption cannot be sustained.
The evaluation in Section 3.2 will address these issues and provide ideas for improved representations of real-valued emissions for the use in spectral algorithms.

Chapter 3

Results and discussion

This chapter will present and discuss the experiments conducted in this work. Section 3.1 characterizes the spectral learning algorithms. Using synthetic experimental data, the algorithms LEARNHKZ and LEARNAHK will be compared with each other. Furthermore, the algorithm LEARNAHK will be compared to the Baum-Welch algorithm. The Baum-Welch algorithm was chosen as an example of a common iterative method for learning hidden Markov models. Section 3.2 then addresses how real-valued emissions can be used for learning parameter estimates using the spectral algorithms.

3.1 Comparison of spectral learning algorithms

3.1.1 Learning synthetic HMMs

The algorithms LEARNHKZ and LEARNAHK were evaluated using synthetic experimental data generated from four different example hidden Markov models. The first two models have a state dimension of $k = 2$ and the second pair of models has a state dimension of $k = 3$. For the latter two models, a larger number of possible emissions was chosen. The parameters of the models used are

$\mathbf{k = 2, d = 3}$	$\mathbf{k = 2, d = 6}$
$T = \begin{pmatrix} 9/10 & 3/10 \\ 1/10 & 7/10 \end{pmatrix}$	$T = \begin{pmatrix} 9/10 & 1/20 \\ 1/10 & 19/20 \end{pmatrix}$
$O = \begin{pmatrix} 1/4 & 8/10 \\ 1/2 & 1/10 \\ 1/4 & 1/10 \end{pmatrix}$	$O = \begin{pmatrix} 1/6 & 7/12 \\ 1/6 & 1/12 \\ 1/6 & 1/12 \\ 1/6 & 1/12 \\ 1/6 & 1/12 \\ 1/6 & 1/12 \end{pmatrix}$
$\vec{\pi} = (4/5, 1/5)^\top$	$\vec{\pi} = (3/4, 1/4)^\top$

and

$\mathbf{k = 3, d = 8}$	$\mathbf{k = 3, d = 10}$
$T = \begin{pmatrix} 8/10 & 1/15 & 1/8 \\ 1/10 & 13/15 & 1/8 \\ 1/10 & 1/15 & 3/4 \end{pmatrix}$	$T = \begin{pmatrix} 8/10 & 1/15 & 1/6 \\ 1/10 & 13/15 & 1/6 \\ 1/10 & 1/15 & 2/3 \end{pmatrix}$
$O = \begin{pmatrix} 3/10 & 1/20 & 1/50 \\ 1/10 & 13/20 & 1/50 \\ 1/10 & 1/20 & 22/50 \\ 1/10 & 1/20 & 22/50 \\ 1/10 & 1/20 & 1/50 \\ 1/10 & 1/20 & 1/50 \\ 1/10 & 1/20 & 1/50 \\ 1/10 & 1/20 & 1/50 \end{pmatrix}$	$O = \begin{pmatrix} 6/15 & 1/20 & 1/50 \\ 1/15 & 11/20 & 1/50 \\ 1/15 & 1/20 & 21/50 \\ 1/15 & 1/20 & 21/50 \\ 1/15 & 1/20 & 1/50 \\ 1/15 & 1/20 & 1/50 \\ 1/15 & 1/20 & 1/50 \\ 1/15 & 1/20 & 1/50 \\ 1/15 & 1/20 & 1/50 \\ 1/15 & 1/20 & 1/50 \end{pmatrix}$
$\vec{\pi} = (1/3, 1/3, 1/3)^\top$	$\vec{\pi} = (1/3, 1/3, 1/3)^\top$

The model for $k = 2$, $d = 6$ was taken from the MathWorks Matlab R2012b documentation. The matrices of all four models were chosen due to their low condition numbers. [1]
From each model, N sequences of triples of categorical emissions were generated, with

$$N = 1000, 2500, 5000, 10000, 25000, 50000, 100000.$$

These sequences of triples were then used as the input for the learning algorithms. For each N, 100 realizations of the sequences were generated.

[1] The condition number of a matrix is the ratio of the largest singular value to the smallest singular value of that matrix.

Both algorithms return empirical estimates for the state transition probability matrix $\hat{T}$ and the emission probability matrix $\hat{O}$. These empirical estimates were compared to the known true values for T and O under the squared Frobenius norm $\|\cdot\|_F^2$. Furthermore, the running times of the algorithms were recorded. The 100 realizations then allowed to calculate average estimation errors and average running times for the different learning instances.
All simulations were implemented using the software Matlab on a personal computer platform. The specifications of the computing environment can be found in Appendix A.2. The source code for the implementations of the algorithms LEARNHKZ and LEARNAHK is available online, see Appendix A.3.

3.1.2 Accuracy and dimensionality

The two spectral algorithms were compared with respect to the accuracy of the estimated model parameters and their running time. Figure 3.1 shows the results for two-state HMMs, and Figure 3.2 shows the results for three-state HMMs.

Two-state HMMs

For the two-state hidden Markov model with three emissions (Figure 3.1a), both algorithms show the same level of accuracy. Furthermore the accuracy significantly increases with the number of triples of observations that used for learning.
When learning from experimental data sampled from the two-state HMM with six emissions, the algorithm LEARNAHK gives more accurate estimates than LEARNHKZ, see Figure 3.1b.
Figures 3.1a and 3.1b show the data on a $\log-\log$ scale. The average slopes of the estimation errors are given in the following table:

	$k=2, d=3$	$k=2, d=6$
LEARNHKZ	-1.08	-0.91
LEARNAHK	-1.06	-1.03

The slopes were obtained using a linear fit on the data over the sample size N, excluding $N = 1000$.
The estimation error of both learning algorithms therefore decreases ap-

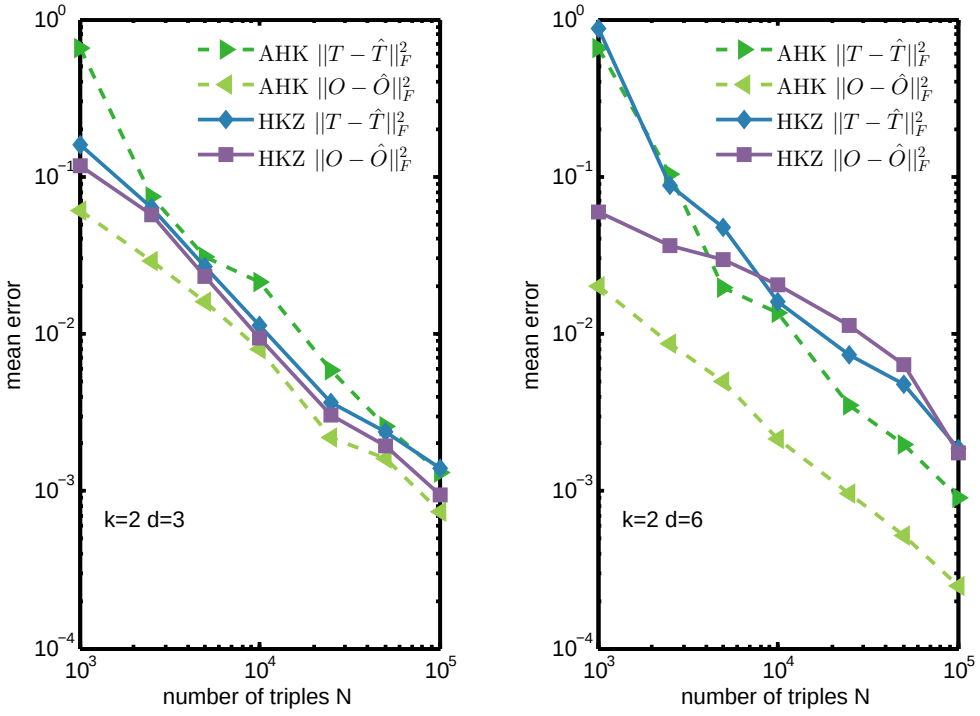

(a) Mean errors $\mathbf{k = 2}$, $\mathbf{d = 3}$. **(b)** Mean errors $\mathbf{k = 2}$, $\mathbf{d = 6}$.

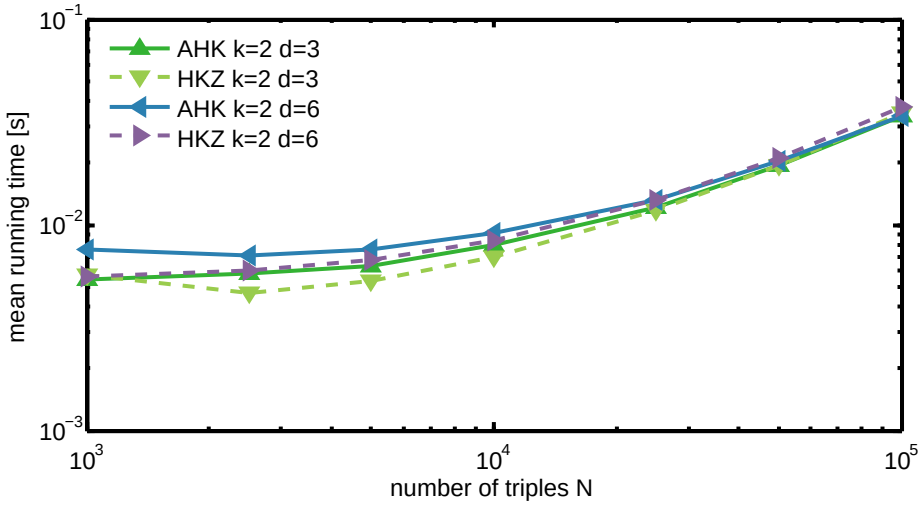

(c) Mean running time for $\mathbf{k = 2}$.

Figure 3.1: Comparison of the accuracy and running time of LEARNHKZ and LEARNAHK for $\mathbf{k = 2}$.

proximately exponentially with the sample size.
Figure 3.1c shows the running time of both algorithms. Also depicted on a $\log - \log$ scale, both algorithms show equal running time, increasing polynomially with the sample size.

Three-state HMMs

In a second step, the algorithms were compared using data from three-state HMMs with eight and ten possible emissions. The results are shown in Figure 3.2.
The estimation errors for learning the three-state eight-emission and three-state ten-emission HMM are depicted in Figures 3.2a and 3.2b, respectively. In contrast to the two-state HMMs with a lower number of possible emissions, the algorithm LEARNHKZ does not show a decreasing estimation error with increasing sample size. The estimation error for learning the transition probability matrix $\hat{T}$ even increases with the sample size. The algorithm LEARNAHK only shows an improving accuracy for $\hat{T}$. For both models and both algorithms, the estimation error of $\hat{O}$ does not decrease with increasing sample size.

This behavior of the learning algorithms is not expected. Ideally, the accuracy should improve for larger sample sizes, just as observed for the two-state HMMs above. The matrices used to generate the observation sequences were chosen to be well-conditioned. However, the matrix decompositions and inversions in the algorithms are prone to instability when the experimentally obtained matrices are ill-conditioned. Especially for matrices with high dimensionality, instability during the numerical operations used in these algorithms is an omnipresent issue.
The decreasing accuracy for the empirical estimate $\hat{T}$ observed for the algorithm LEARNHKZ could be due to another issue: the method used to obtain the parameter estimates does neither guarantee probabilities in the range $[0, 1]$, nor does it observe the scaling of the columns of the transition probability matrix that is estimated [Hsu *et al.* 2012]. Both these issues are likely to be contributing to this unusual behavior.
The running time for both algorithms, depicted in Figure 3.2c, is approximately the same. This is consistent with the results for two-state HMMs.

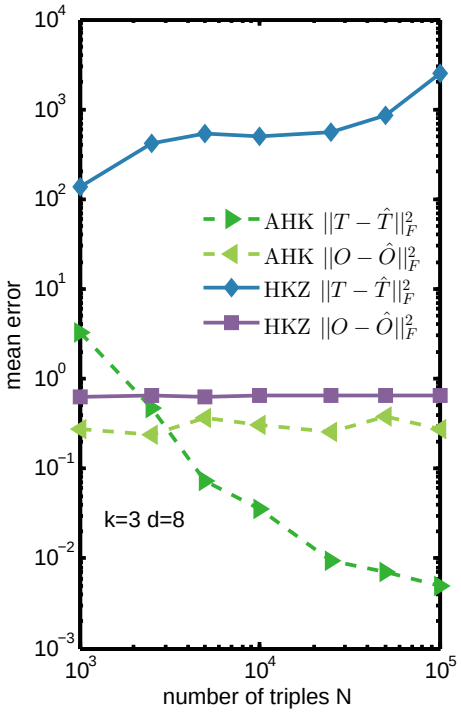

(a) Mean errors $\mathbf{k = 3}$, $\mathbf{d = 8}$.

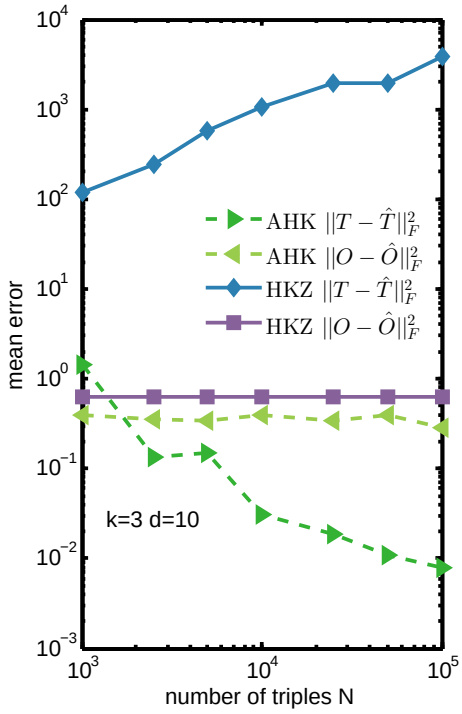

(b) Mean errors $\mathbf{k = 3}$, $\mathbf{d = 10}$.

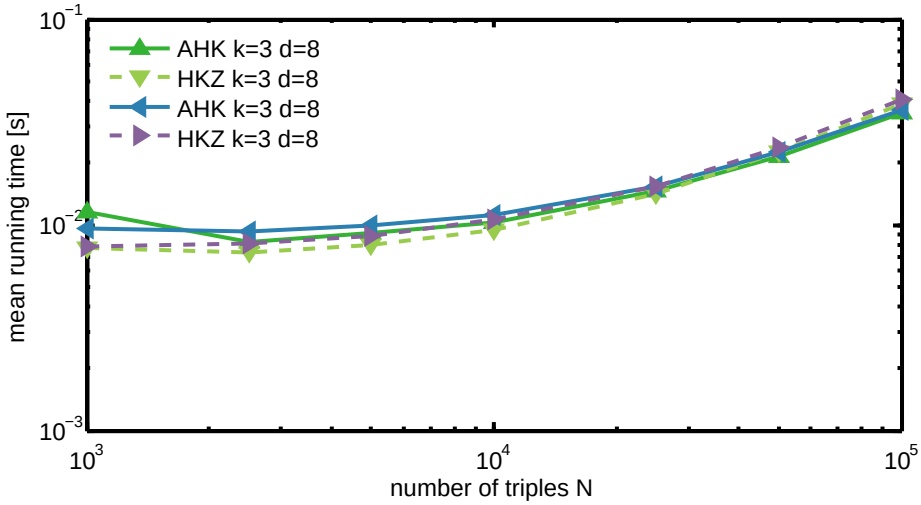

(c) Mean running time for $\mathbf{k = 3}$.

Figure 3.2: Comparison of the accuracy and running time of LEARNHKZ and LEARNAHK for $\mathbf{k = 3}$.

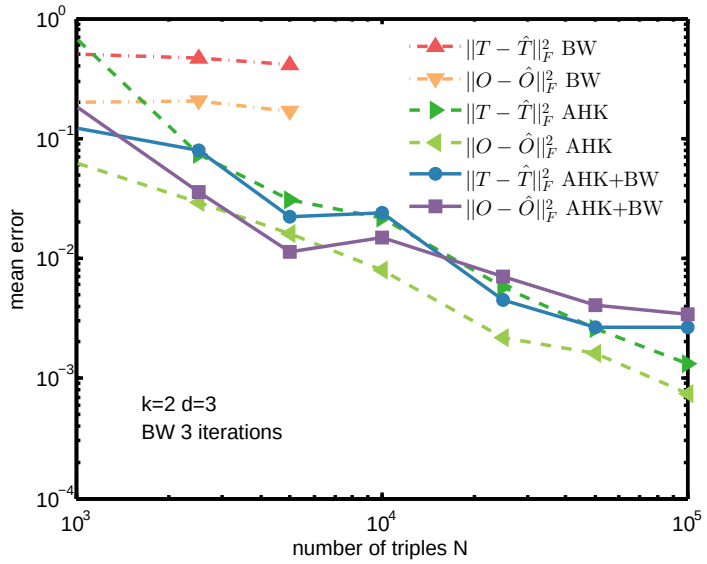

(a) Mean errors for $\hat{O}$ and $\hat{T}$

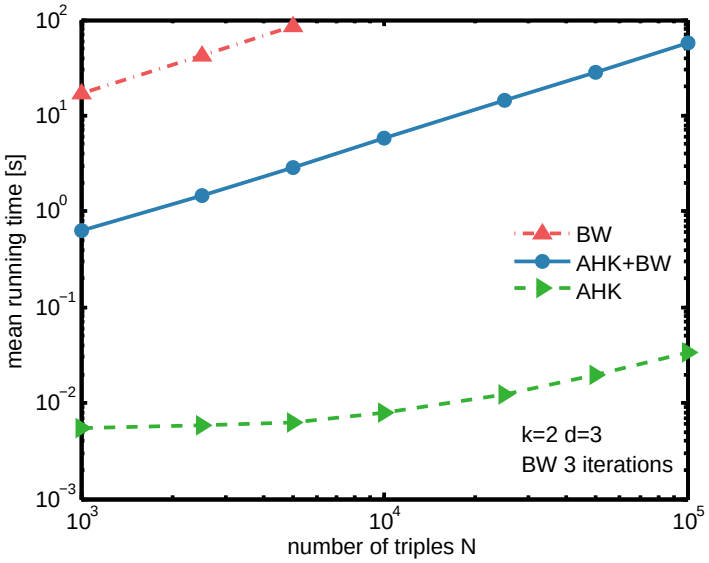

(b) Mean running times.

Figure 3.3: Comparison of the algorithm LEARNAHK (indicated by *AHK*) with the Baum-Welch algorithm, both with a random initialization (*BW*), and using initial values from the LEARNAHK algorithm (*AHK+BW*). Average errors and running times across 100 realizations for *AHK* and 10 realizations for *BW* and *AHK+BW*

3.1.3 Comparison with the Baum-Welch algorithm

In a next step, the algorithm LEARNAHK was compared to the Baum-Welch algorithm. For the simulations, the implementation of the Baum-Welch algorithm as provided by the Matlab Statistics Toolbox was used. In this experiment, only sequences sampled from the two-state three-emission ($k = 2$, $d = 3$) hidden Markov model specified earlier were considered. Using the same model experimental data as before, the Baum-Welch (BW) algorithm was run for learning the parameter estimates of the HMM.
The spectral algorithms do not require initial estimates for the quantities to be learned. With the BW algorithm being an iterative method, however, initial estimates are required for initialization. Therefore, two different scenarios were considered: the BW algorithm with a random initialization, and the BW algorithm initialized with the parameters learned by LEARNAHK.
The results are shown in Figure 3.3. For the BW algorithm with random initialization, only sample sizes $N = 1000, 2500, 5000$ were considered, due to the long running time. In both instances, the BW algorithm was run for three iterations. The results for the BW algorithm are averages over 10 realizations of input sequences.
Figure 3.3b shows the running time of the different learning instances. It can be seen that already for the smallest sample size, $N = 1000$, the BW algorithm is much slower than the spectral algorithm. The difference in running time is more than three orders of magnitude. Compared to the BW algorithm initialized with learning results from LEARNAHK, the spectral algorithm by itself is still two orders of magnitude faster for the smallest sample size. While the running time of the spectral algorithm only increases polynomially, it increases exponentially for the BW algorithm. This can be seen by the straight lines for the running times of the BW algorithm on the $\log - \log$ scale. The differences in running time will therefore become even larger for increasing sample size.
Figure 3.3a shows the accuracy of the different learning instances. For three iterations, the accuracy of the BW algorithm with random initialization does not significantly increase for a larger sample size. Furthermore, the estimated parameters are not very accurate. For the BW algorithm initialized with learning results from the spectral algorithm, three iterations do not result in more accurate parameter estimates.
The spectral algorithm is significantly faster than the Baum-Welch algorithm. Already for small sample sizes, this advantage in speed is by many orders of magnitude. Only having been run for three iterations, the parameter estimates resulting from the BW algorithm are not more accurate than

the results of the spectral algorithm. A further improvement of the solution would require many more iterations, further increasing the running time. A work by [Kontorovich *et al.* 2013] addresses the convergence of the BW algorithm with respect to the number of iterations.

3.2 Real-valued emissions

In this section, the accuracy of parameter estimates obtained from sequences of real-valued emissions is evaluated. In contrast to the analysis in Section 3.1, sequences of real-valued emissions cannot be used as an input for the learning algorithms right away. It is therefore important to use a suitable method to prepare the data for the algorithms first.
One possible procedure, binning the emissions according to quantiles, was introduced in Section 2.3. In order to be able to quantify the accuracy of the parameter estimates learned from the data, synthetic sequences of real-valued observations were used.

3.2.1 Generative process for synthetic experimental data

Sequences of categorical emissions were generated from the hidden Markov models specified in the beginning of Section 3.1. Only the two-state HMMs were considered. These sequences were then used to generate sequences of real valued emissions, resulting in our model experimental data.
Let there be a sequence of categorical emissions. The emissions are then used as the means for the generation of normally distributed random numbers with a fixed standard deviation σ. This results in a sequence of real-valued emissions. This generative process is depicted in Figure 3.4.

Learning from these sequences should therefore result in parameter estimates for the emission probability matrices and state transition matrices of the HMMs considered in this experiment.

3.2.2 Simple binning approach

The model experimental data used to evaluate the simple binning approach was generated as specified by the generative process above. The model two-state HMMs ($k = 2$, $d = 3$ and $k = 2$, $d = 6$) are given in the beginning of Section 3.1. Series of emissions generated using two different standard deviations, $\sigma = 0.1$ and $\sigma = 0.25$, were considered.
The following procedure was used to assess the simple binning approach:

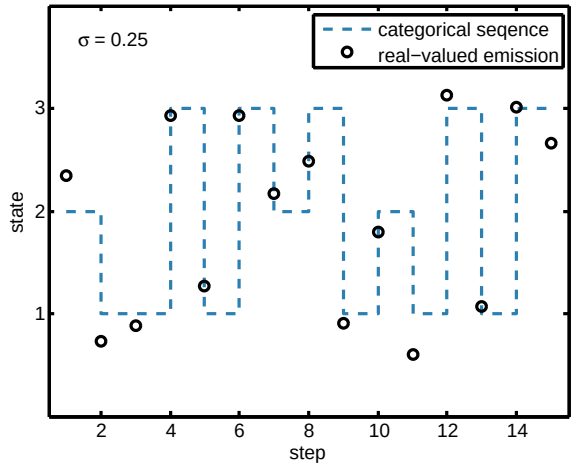

Figure 3.4: The generative process for a sequences of real-valued emissions. The categorical emissions (blue dashed line) were used as the mean values to generate a sequence of normally distributed random numbers with fixed standard deviation (here $\sigma = 0.25$).

1. Given a sequence of real-valued emissions, use SIMPLEBINNING to bin the emissions, creating a sequence of categorical emissions.

2. Use this sequence of categorical emissions for learning the HMM model parameters with a spectral algorithm.

3. Compare the resulting estimate of the emission probability matrix $\hat{O}$ and state transition probability matrix $\hat{T}$ to the known true values.

Only the algorithm LEARNAHK was used for this analysis, since it was established in Section 3.1 that this algorithm is the more accurate one of the two algorithms considered in this work.

The results for the simple binning approach are shown in Figure 3.5. The accuracy of the parameter estimates as obtained by the procedure specified above was compared against the sample size.
For learning estimates of the state transition probability matrix $\hat{T}$, both datasets result in an equally accurate learning, with the estimation error decreasing with increasing sample size. This corresponds to the same behavior as observed for the categorical case.

Considering the case for an estimate of the emission probability matrix $\hat{O}$, we see a different picture. Here an increased sample size does not result in a significant decrease in the estimation error. We can compare this to the results in Section 3.1, where the emission probability matrix was

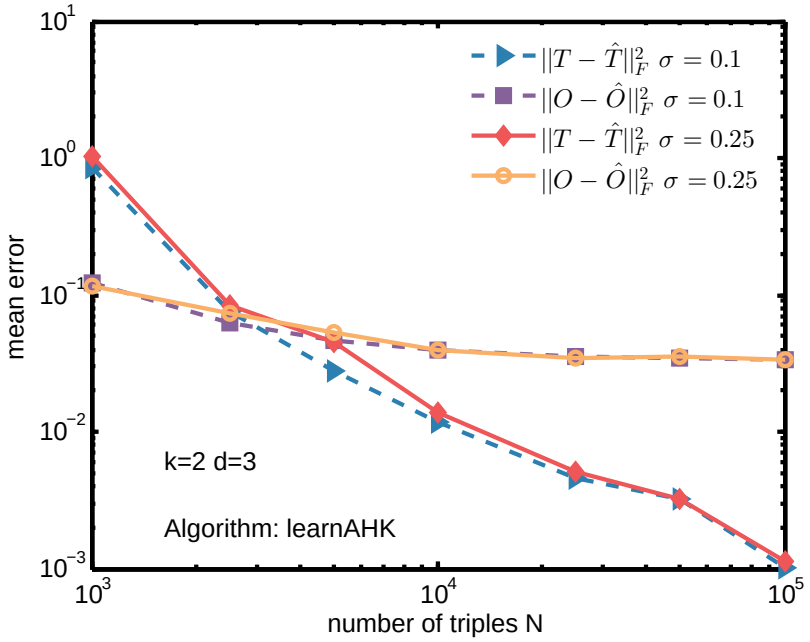

(a) $k = 2, d = 3$

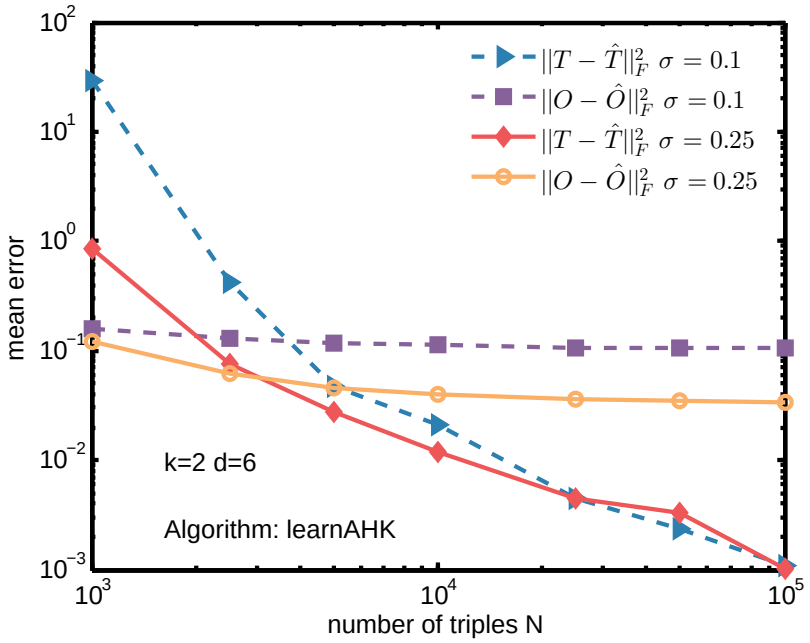

(b) $k = 2, d = 6$

Figure 3.5: Accuracy of the parameter estimates using the simple binning approach. Two-state HMMs with three and six possible emissions, respectively.

properly learned for categorical emissions. The real-valued emissions used in this analysis were based on the same model parameters. The inability to learn the estimate $\hat{O}$ is very likely to be due to the problem associated with equally sized bins when calculating the quantiles, as discussed in 2.3.1.
The problems associated with the simple binning approach should not lead to the idea of binning by quantiles to be discarded though.

3.2.3 Alternative approaches

One possible solution to the problems associated with the simple binning approach is to bin on a finer scale. Instead of just binning into as many bins as the expected dimension of the emission probability matrix, one can consider using a larger number of bins.
Each time a real-valued observation is binned, information gets lost. The strength of the spectral methods lies partly in the fact that they can readily handle high observation spaces. This can be utilized with respect to the number of bins. When considering all sequences of real-valued emissions available, a larger number of quantiles can be calculated, resulting in a larger number of bins. The binning should ideally be finer than the smallest probability to be expected for a possible emission. Binning on a finer scale and retaining the information about the bin bounds leads to less information being lost.
When learning from time series binned on a finer scale, the spectral algorithms will return emission probability matrices of larger dimension than initially desired. Together with the bin bounds, these matrices can be used to recover probability distribution functions.
The result would be a mixture of probability distribution functions. Learning from mixtures of distributions is an active field of research and beyond the scope of this work. Among various available learning methods, spectral algorithms have also been developed for mixtures of distributions, see for example [Achlioptas & Mcsherry 2005] and [Dasgupta 1999].
The algorithm by [Hsu *et al.* 2012] is only formulated for sequences of discrete observations. The algorithm by [Anandkumar *et al.* 2012b], however, also allows multivariate continuous observations. For this algorithm, categorical observations are represented by a basis vector in the d-dimensional standard coordinate system (see 2.2.1). It would also be possible to use a sequence of d-dimensional probability vectors as an input for this algorithm.
Instead of binning real-valued observations, where the challenge lies in finding appropriate bin bounds, one could therefore also try to find a rep-

resentation of a scalar observation by a d-dimensional probability vector. The resulting sequence of vectors could then be used as the input for the learning algorithm LEARNAHK.

Chapter 4

Conclusion

Spectral methods for learning hidden Markov models provide an appealing alternative to existing learning algorithms based on local search heuristics. The computational and sample complexity for existing search heuristics is oftentimes exponential, while it is only polynomial for the spectral methods considered in this work. Furthermore, spectral methods provide global estimates and are not prone to local optima. Both advantages become especially important when working with large observation spaces and large numbers of samples.

The synthetic observation sequences used in this work allowed to compare the obtained parameter estimates for example HMMs to their true values. It was established that the algorithm by [Anandkumar *et al.* 2012b] provides better learning results than the one proposed by [Hsu *et al.* 2012]. Both algorithms handle large observation spaces computationally efficiently.

For the state transition probability matrices, only the algorithm by [Anandkumar *et al.* 2012b] learned reasonable parameter estimates for all four example HMMs considered. The algorithm by [Hsu *et al.* 2012] did not obtain reasonable estimates for two of the example HMMs. It should be noted again, however, that the algorithm by [Hsu *et al.* 2012] in its original form does not provide estimates for the HMM parameter matrices, and an additional computation step had to be implemented. The problems encountered for some of the example HMMs using the algorithm LEARNHKZ are consistent with instability concerns regarding this additional computation step by the authors in the original work.

Both algorithms failed to learn estimates for the emission probability matrices for two of the example HMMs, even for a large number of samples. These two HMMs were the ones with the largest observation spaces. Circumstances that are likely to have contributed to this are: instabilities encountered in numerical singular value and eigenvalue decompositions, the

scaling of columns, and learned probabilities outside the interval $[0, 1]$.
Comparing the algorithm by [Anandkumar *et al.* 2012b] to the Baum-Welch algorithm, we found advantages in running time and accuracy of estimated quantities. Already for small sample sizes, the spectral algorithm was several orders of magnitude faster than the iterative method. The computational and sample complexity of polynomial order lets this advantage become even greater for larger sample sizes. This is even more significant when considering that the Baum-Welch algorithm was only run for few iterations. More iterations, which are generally needed for more accurate results, would increase the running time even further. Together with the spectral methods being free of local optima, this highlights major advantages over iterative methods.
The representation of sequences of real-valued observations to be used in the spectral algorithms proved to be challenging. The initial simple binning approach turned out to be not suitable for learning an estimate for the emission probability matrix. The state transition probability matrix was learned comparably well. This simple approach did not make full use of the advantages provided by the spectral methods though. The proposed idea to use a finer binning is only one possible approach for a better representation of real-valued emissions. However, this approach would require additional work after parameter estimates are obtained with the spectral algorithms.

Outlook

Spectral algorithms for learning hidden Markov models emerged only rather recently, and the two algorithms presented in this report are important examples of this approach. The tensor decomposition proposed by [Anandkumar *et al.* 2012a] provides a method that might be less susceptible to instability. A possible future work would be to formulate an explicit learning algorithm for hidden Markov models based on this method, and compare it to the spectral methods considered in this work.
Regarding real-valued observations, further investigation is needed to find robust and reasonable representations for the use with spectral algorithms. Binning on a finer scale or representing real-valued observations by probability vectors are both two possible starting points for future considerations.

Appendix A

Appendix

A.1 Proofs

Proposition 1

Proof. Observing the conditional independence and using Bayes' theorem we find that

$$\begin{aligned}\Pr\left[h_1 = i | h_2 = j\right] &= \frac{\Pr\left[h_2 = j | h_1 = i\right] \cdot \Pr\left[h_1 = i\right]}{\Pr\left[h_2 = j\right]} \\ &= \frac{T_{ji}\pi_i}{\left(T\vec{\pi}\right)_j} = \vec{e}_i \operatorname{diag}\left(\vec{\pi}\right) T^\top \operatorname{diag}\left(T\vec{\pi}\right)^{-1} \vec{e}_j.\end{aligned}$$

With definition 2.2.1 we therefore obtain

$$M_1\vec{e}_j = \mathbb{E}\left[\vec{x}_1 | h_2 = j\right] = O\mathbb{E}\left[\vec{e}_{h_1} | h_2 = j\right] = O \operatorname{diag}\left(\vec{\pi}\right) T^\top \operatorname{diag}\left(T\vec{\pi}\right)^{-1} \vec{e}_j.$$

With the definitions of the observation probability matrix O and the state transition probability matrix T from section 1.4 we find

$$M_2\vec{e}_j = \mathbb{E}\left[\vec{x}_2 | h_2 = j\right] = O\vec{e}_j$$

and

$$M_3\vec{e}_j = \mathbb{E}\left[\vec{x}_3 | h_2 = j\right] = O\mathbb{E}\left[\vec{e}_{h_3} | h_2 = j\right] = OT\vec{e}_j.$$

□

Lemma 1

Proof. Observing conditional independence, we have

$$\begin{aligned} P_{3,1} &= \mathbb{E}\left[\mathbb{E}\left[\vec{x}_3 \otimes \vec{x}_1 | h\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\vec{x}_3 | h\right] \otimes \mathbb{E}\left[\vec{x}_1 | h\right]\right] \\ &= \mathbb{E}\left[(M_3 \vec{e}_h) \otimes (M_1 \vec{e}_h)\right] \\ &= M_3 \left(\sum_{t=1}^{k} w_t \vec{e}_t \otimes \vec{e}_t\right) M_1^\top \\ &= M_3 \mathrm{diag}(\vec{w}) M_1^\top \end{aligned}$$

and

$$\begin{aligned} P_{3,1,2} &= \mathbb{E}\left[\mathbb{E}\left[(\vec{x}_3 \otimes \vec{x}_1)\left\langle \vec{\eta}, \vec{x}_2 \right\rangle | h\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\vec{x}_3 | h\right] \otimes \mathbb{E}\left[\vec{x}_1 | h\right]\left\langle \vec{\eta}, \mathbb{E}\left[\vec{x}_2 | h\right]\right\rangle\right] \\ &= \mathbb{E}\left[(M_3 \vec{e}_h) \otimes (M_1 \vec{e}_h)\left\langle \vec{\eta}, M_2 \vec{e}_h \right\rangle\right] \\ &= M_3 \left(\sum_{t=1}^{k} w_t \vec{e}_t \otimes \vec{e}_t \left\langle \vec{\eta}, M_2 \vec{e}_h \right\rangle\right) M_1^\top \\ &= M_3 \,\mathrm{diag}(M_2^\top \vec{\eta})\,\mathrm{diag}(\vec{w})\, M_1^\top . \end{aligned}$$

□

Lemma 2

Proof. Using Lemma 1,

$$U_3^\top P_{3,1}(\vec{\eta}) U_1 = \left(U_3^\top M_3\right) \mathrm{diag}(\vec{w}) \left(M_1^\top U_1\right)$$

is invertible by the assumption on U_t and Condition 2.2.1. By using the relations from Lemma 1 we obtain

$$\begin{aligned} B_{3,1,2}(\vec{\eta}) &= \left(U_3^\top P_{3,1,2}(\vec{\eta}) U_1\right)\left(U_3^\top P_{3,1} U_1\right)^{-1} \\ &= \left(U_3^\top M_3 \,\mathrm{diag}(M_2^\top \vec{\eta})\,\mathrm{diag}(\vec{w})\, M_1^\top U_1\right)\left(U_3^\top P_{3,1} U_1\right)^{-1} \\ &= \left(U_3^\top M_3\right)\mathrm{diag}(M_2^\top \vec{\eta})\left(U_3^\top M_3\right)^{-1}\left(U_3^\top M_3\right)\mathrm{diag}(\vec{w})\left(M_1^\top U_1\right)\left(U_3^\top P_{3,1} U_1\right)^{-1} \\ &= \left(U_3^\top M_3\right)\mathrm{diag}(M_2^\top \vec{\eta})\left(U_3^\top M_3\right)^{-1}\left(U_3^\top P_{3,1} U_1\right)\left(U_3^\top P_{3,1} U_1\right)^{-1} \\ &= \left(U_3^\top M_3\right)\mathrm{diag}(M_2^\top \vec{\eta})\left(U_3^\top M_3\right)^{-1} . \end{aligned}$$

□

Lemma 3

Proof. From Equation (2.5) and by setting $\vec{\eta} = U_2\vec{\theta}_i$ we find $\forall i \in [k]$

$$\begin{aligned}\left(U_3^\top M_3\right)^{-1} B_{3,1,2}(U_2\vec{\theta}_i)\left(U_3^\top M_3\right) &= \operatorname{diag}(M_2^\top U_2\vec{\theta}_i)\\ &= \operatorname{diag}(\langle\vec{\theta}_i, U_2^\top M_2\vec{e}_1\rangle, \langle\vec{\theta}_i, U_2^\top M_2\vec{e}_2\rangle, \dots \langle\vec{\theta}_i, U_2^\top M_2\vec{e}_k\rangle\\ &= \operatorname{diag}\left(\lambda_{i,1}, \lambda_{i,2}, \dots, \lambda_{i,k}\right).\end{aligned}$$

Using the definition of $\Theta \in \mathbb{R}^{k\times k}$, we obtain for $L \in \mathbb{R}^{k\times k}$

$$\begin{aligned}L &= \begin{pmatrix}\lambda_{1,1} & \lambda_{1,2} & \cdots & \lambda_{1,k}\\ \lambda_{2,1} & \lambda_{2,2} & \cdots & \lambda_{2,k}\\ \vdots & \vdots & \ddots & \vdots\\ \lambda_{k,1} & \lambda_{k,2} & \cdots & \lambda_{k,k}\end{pmatrix}\\ &= \begin{pmatrix}\langle\vec{\theta}_1, U_2^\top M_2\vec{e}_1\rangle & \langle\vec{\theta}_1, U_2^\top M_2\vec{e}_2\rangle & \cdots & \langle\vec{\theta}_1, U_2^\top M_2\vec{e}_k\rangle\\ \langle\vec{\theta}_2, U_2^\top M_2\vec{e}_1\rangle & \langle\vec{\theta}_2, U_2^\top M_2\vec{e}_2\rangle & \cdots & \langle\vec{\theta}_2, U_2^\top M_2\vec{e}_k\rangle\\ \vdots & \vdots & \ddots & \vdots\\ \langle\vec{\theta}_k, U_2^\top M_2\vec{e}_1\rangle & \langle\vec{\theta}_k, U_2^\top M_2\vec{e}_2\rangle & \cdots & \langle\vec{\theta}_k, U_2^\top M_2\vec{e}_k\rangle\end{pmatrix}\\ &= \Theta U_2^\top M_2.\end{aligned}$$

□

A.2 Specifications of computing environment

Computations were conducted on a personal computer.

Relevant specifications:

- Intel Core i5-540M
- 4.00 GB RAM
- 160 GB Intel solid state drive
- Microsoft Windows 8.1 Professional 64-bit
- MathWorks Matlab R2012b 64-bit

A.3 Source code

The algorithms LEARNHKZ and LEARNAHK were implemented in Matlab. The source code, along with the binning routine used in this work, is available at https://github.com/cmgithub/spectral.

Bibliography

[Achlioptas & Mcsherry 2005] Dimitris Achlioptas and Frank Mcsherry. *On spectral learning of mixtures of distributions*. Learning Theory, pages 458–469, 2005. (Cited on page 36.)

[Anandkumar *et al.* 2012a] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade and Matus Telgarsky. *Tensor decompositions for learning latent variable models*. arXiv preprint arXiv:1210.7559, pages 1–55, October 2012. (Cited on pages 3, 14 and 40.)

[Anandkumar *et al.* 2012b] Animashree Anandkumar, Daniel Hsu and Sham M. Kakade. *A Method of Moments for Mixture Models and Hidden Markov Models*. arXiv preprint arXiv:1203.0683, pages 1–31, March 2012. (Cited on pages v, 2, 3, 11, 13, 14, 16, 17, 18, 19, 21, 36, 39 and 40.)

[Baum & Petrie 1966] Leonard E Baum and Ted Petrie. *Statistical Inference for Probabilistic Functions of Finite State Markov Chains*. The Annals of Mathematical Statistics, vol. 37, no. 6, pages 1554–1563, 1966. (Cited on page 4.)

[Baum *et al.* 1970] Leonard E Baum, Ted Petrie, George Soules and Norman Weiss. *A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains*. The annals of mathematical statistics, vol. 41, no. 1, pages 164–171, 1970. (Cited on page 1.)

[Bilmes 1998] Jeff A Bilmes. *A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models*. International Computer Science Institute, vol. 4, no. 510, page 126, 1998. (Cited on page 1.)

[Bishop 2006] Christopher M Bishop. Pattern Recognition and Machine Learning. Springer, New York, New York, USA, 2006. (Cited on pages 2, 4 and 8.)

[Chang 1996] Joseph T Chang. *Full reconstruction of Markov models on evolutionary trees: Identifiability and consistency*. Mathematical Biosciences, vol. 137, no. 1, pages 51–73, October 1996. (Cited on pages 1 and 3.)

[Dasgupta 1999] Sanjoy Dasgupta. *Learning mixtures of Gaussians*. 40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039), pages 634–644, 1999. (Cited on page 36.)

[Dempster *et al.* 1977] Arthur P Dempster, Nan M Laird and Donald B Rubin. *Maximum likelihood from incomplete data via the EM algorithm*. Journal of the Royal Statistical Society, vol. 39, no. 1, pages 1–38, 1977. (Cited on page 1.)

[Hsu *et al.* 2012] Daniel Hsu, Sham M. Kakade and Tong Zhang. *A spectral algorithm for learning Hidden Markov Models*. Journal of Computer and System Sciences, vol. 78, no. 5, pages 1460–1480, September 2012. (Cited on pages v, 1, 2, 3, 11, 12, 13, 18, 29, 36 and 39.)

[Jaeger 2000] Herbert Jaeger. *Observable operator models for discrete stochastic time series*. Neural computation, vol. 12, no. 6, pages 1371–98, June 2000. (Cited on page 18.)

[Kontorovich *et al.* 2013] Aryeh Kontorovich, B Nadler and Roi Weiss. *On learning parametric-output HMMs*. arXiv preprint arXiv:1302.6009, vol. 28, 2013. (Cited on page 33.)

[Krogh *et al.* 2001] Anders Krogh, Björn Larsson, Gunnar von Heijne and Erik L L Sonnhammer. *Predicting transmembrane protein topology with a hidden Markov model: application to complete genomes*. Journal of molecular biology, vol. 305, no. 3, pages 567–80, January 2001. (Cited on page 1.)

[Manning & Schütze 1999] Christopher D Manning and Hinrich Schütze. Foundations of statistical natural language processing. MIT Press, Cambridge, MA, USA, 1999. (Cited on page 1.)

[Mossel & Roch 2005] Elchanan Mossel and Sébastien Roch. *Learning nonsingular phylogenies and hidden Markov models*. Proceedings of the thirty-seventh annual ACM symposium on Theory of computing - STOC '05, page 366, 2005. (Cited on page 3.)

[Pearson 1894] Karl Pearson. *Contributions to the mathematical theory of evolution*. Philosophical Transactions of the Royal Society of London. A, vol. 185, no. 1894, pages 71–110, 1894. (Cited on page 1.)

[Rabiner 1989] Lawrence R Rabiner. *A tutorial on hidden Markov models and selected applications in speech recognition*. Proceedings of the IEEE, vol. 77, no. 2, pages 257–286, 1989. (Cited on pages 1 and 8.)

[Redner & Walker 1984] Richard A Redner and Homer F Walker. *Mixture Densities, Maximum Likelihood and the EM Algorithm*. SIAM Review, vol. 26, no. 2, pages 195–239, April 1984. (Cited on page 1.)

[Siddiqi *et al.* 2009] Sajid M. Siddiqi, Byron Boots and Geoffrey J. Gordon. *Reduced-Rank Hidden Markov Models*. arXiv preprint, vol. arXiv:0910, October 2009. (Cited on page 3.)